EXHIBIT G

EXHIBIT 23

EXHIBIT FILED UNDER SEAL

UBER_JCCP_MDL_003231342

Metadata

#Author	actuaryzhang@uber.com	SEMANTIC
#Date Modified	01/28/2019	SEMANTIC
#DateCreated	03/12/2016	SEMANTIC
#Title	Bouncer v3	SEMANTIC
Account	brooke.anderson@uber.com;	SEMANTIC
All Custodians	Anderson, Brooke;Breeden, Tracey;Brown, Greg;Cardenas, Philip;Chang, Frank;Fuldner, Gus;Kaiser, Roger;Kansal, Sachin;Lake, Carley;McDonald, Katy;Parker, Kate;Sheridan, Danielle;Silver, Nick;Sullivan, Joe;	SEMANTIC
All Paths	Anderson, Brooke: \20240905\Anderson, Brooke\Drive\EDISCO-25640_DR_25.zip; Anderson, Brooke: \20240905\Anderson, Brooke\Drive\EDISCO-25640_DR_25.zip; Anderson, Brooke: \MassTort_Category4_DRIVE\MassTort_Category4_DRIVE_6.zip; Anderson, Brooke: \MassTort_Category4_DRIVE\MassTort_Category4_DRIVE_6.zip; Breeden, Tracey: \JCCP_DRIVE002_002\JCCP_DRIVE002_002_18.zip; Breeden, Tracey: \JCCP_DRIVE002_002\JCCP_DRIVE002_002_18.zip; Breeden, Tracey: \MassTort_Category2_DRIVE\MassTort_Category2_DRIVE_9.zip; Breeden, Tracey: \MassTort_Category2_DRIVE\MassTort_Category2_DRIVE_9.zip; Breeden, Tracey: \MassTort_Category4_DRIVE\MassTort_Category4_DRIVE_6.zip; Breeden, Tracey: \MassTort_Category4_DRIVE\MassTort_Category4_DRIVE_6.zip; Brown, Greg: \Drive_003_002_Drive003_001\JCCP_Drive003_001\JCCP_Drive003_001_18.zip; Brown, Greg: \Drive_003_002_Drive003_001\JCCP_Drive003_001\JCCP_Drive003_001_18.zip; Cardenas, Philip: \Drive_003_002_Drive003_001\JCCP_DRIVE003_002\JCCP_DRIVE003_002_1.zip; Cardenas, Philip: \Drive_003_002_Drive003_001\JCCP_DRIVE003_002\JCCP_DRIVE003_002_1.zip; Chang, Frank: \EDISCO-25937_frank@uber.com\EDISCO-25937_frank@uber.com_82.zip; Chang, Frank: \EDISCO-25937_frank@uber.com\EDISCO-25937_frank@uber.com_82.zip; Fuldner, Gus: \JCCP-EDISCO-23800_2019JanMarch\JCCP-EDISCO-23800_2019JanMarch_216.zip; Fuldner, Gus: \JCCP-EDISCO-23800_2019JanMarch\JCCP-EDISCO-23800_2019JanMarch_216.zip; Kaiser, Roger: \JCCP_DRIVE005\JCCP_Drive005_120.zip; Kaiser, Roger: \JCCP_DRIVE005\JCCP_Drive005_120.zip; Kaiser, Roger: \MassTort_Category2_DRIVE\MassTort_Category2_DRIVE_9.zip; Kaiser, Roger: \MassTort_Category2_DRIVE\MassTort_Category2_DRIVE_9.zip; Kansal, Sachin: \EDISCO-25937_sachin.kansal@uber.com\EDISCO-25937_sachin.kansal@uber.com2_65.zip; Kansal, Sachin: \EDISCO-25937_sachin.kansal@uber.com\EDISCO-25937_sachin.kansal@uber.com2_65.zip; Lake, Carley: \EDISCO-24394_Drive\EDISCO-24394_Drive_102.zip; Lake, Carley: \EDISCO-24394_Drive\EDISCO-24394_Drive_102.zip; McDonald, Katy: \JCCP_DRIVE005\JCCP_Drive005_120.zip; McDonald, Katy: \JCCP_DRIVE005\JCCP_Drive005_120.zip; McDonald, Katy: \MassTort_Category2_DRIVE\MassTort_Category2_DRIVE_9.zip; McDonald, Katy: \MassTort_Category2_DRIVE\MassTort_Category2_DRIVE_9.zip; McDonald, Katy: \MassTort_Category4_DRIVE\MassTort_Category4_DRIVE_6.zip; McDonald, Katy: \MassTort_Category4_DRIVE\MassTort_Category4_DRIVE_6.zip; Parker, Kate: \JCCP_DRIVE006\JCCP_DRIVE006_35.zip; Parker, Kate: \JCCP_DRIVE006\JCCP_DRIVE006_35.zip; Sheridan, Danielle: \JCCP-EDISCO-23800_2019JanMarch\JCCP-EDISCO-23800_2019JanMarch_216.zip; Sheridan, Danielle: \JCCP-EDISCO-23800_2019JanMarch\JCCP-EDISCO-23800_2019JanMarch_216.zip; Silver, Nick: \EDISCO-24394_Drive\EDISCO-24394_Drive_102.zip; Silver, Nick: \EDISCO-24394_Drive\EDISCO-24394_Drive_102.zip; Silver, Nick: \MassTort_Category2_DRIVE\MassTort_Category2_DRIVE_9.zip; Silver, Nick: \MassTort_Category2_DRIVE\MassTort_Category2_DRIVE_9.zip; Sullivan, Joe: \EDISCO-25937_joesullivan@uber.com\EDISCO-25937_joesullivan@uber.com_11.zip; Sullivan, Joe: \EDISCO-25937_joesullivan@uber.com\EDISCO-25937_joesullivan@uber.com_11.zip	SEMANTIC
Application	Microsoft 2007 Word Document	SEMANTIC
Attachment Names	word	SEMANTIC
Begin Family	UBER_JCCP_MDL_003231342	SEMANTIC
Collaborators	alexclatterbuck@uber.com; jross@uber.com; bhav@uber.com; mananp@uber.com; butorac@uber.com; brianz@uber.com; dmitriy@uber.com; xin.luo@uber.com; kmcdonald@uber.com; brian.tan@uber.com; doyle@uber.com; rmawas@uber.com; ncf@uber.com; dhruv@uber.com; kchiu@uber.com; sjeon@uber.com; s.team-disabled@uber.com; UNKNOWN; safety-program@uber.com; uber.com	SEMANTIC
Confidentiality	Confidential	SEMANTIC
Date Created	03/12/2016 11:33 pm	SEMANTIC
Date Modified	01/28/2019 10:32 pm	SEMANTIC
DocID	1y4jv644kF23DXdWrpNaVqiXpVgft1vxytSss5cPHUgY	SEMANTIC
Document Type	Electronic File	SEMANTIC
End Family	UBER_JCCP_MDL_003231369	SEMANTIC

EXHIBIT myb
2811
4/16/25 Wong
PENGAD 800-631-6989

File Path	\20240905\Anderson, Brooke\Drive\EDISCO-25640_DR_25.zip	SEMANTIC
File Size	1942794	SEMANTIC
Filename	Bouncer v3_1y4jv644kF23DXdWrpNaVqiXpVgft1vxytSss5cPHUgY.docx	SEMANTIC
GoogleDocumentType	DOCUMENT	SEMANTIC
Hash Value	3552253cd3bcb66a7bdda30abee7192d	SEMANTIC
Hidden Content	Yes;	SEMANTIC
ILS All Bates	UBER_JCCP_MDL_003231342;UBER_JCCP_MDL_003231343;UBER_JCCP_MDL_003231344;UBER_JCCP_MDL_003231345;UBER_JCCP_MDL_003231346;UBER_JCCP_MDL_003231347;UBER_JCCP_MDL_003231348;UBER_JCCP_MDL_003231349;UBER_JCCP_MDL_003231350;UBER_JCCP_MDL_003231351;UBER_JCCP_MDL_003231352;UBER_JCCP_MDL_003231353;UBER_JCCP_MDL_003231354;UBER_JCCP_MDL_003231355;UBER_JCCP_MDL_003231356;UBER_JCCP_MDL_003231357;UBER_JCCP_MDL_003231358;UBER_JCCP_MDL_003231359;UBER_JCCP_MDL_003231360;UBER_JCCP_MDL_003231361;UBER_JCCP_MDL_003231362;UBER_JCCP_MDL_003231363;UBER_JCCP_MDL_003231364;UBER_JCCP_MDL_003231365;UBER_JCCP_MDL_003231366;UBER_JCCP_MDL_003231367;UBER_JCCP_MDL_003231368;UBER_JCCP_MDL_003231369	SEMANTIC
ILS Document Date	01/28/2019	SEMANTIC
ILS Prod Date	3/3/2025	SEMANTIC
ILS Prod Vol	JCCP_MDL106	SEMANTIC
LINKSOURCEBEGBATES	UBER_JCCP_MDL_001728489; UBER_JCCP_MDL_002221404; UBER_JCCP_MDL_002345578; UBER_JCCP_MDL_002345602	SEMANTIC
Other Custodians	Breeden, Tracey;Brown, Greg;Cardenas, Philip;Fuldner, Gus;Kaiser, Roger;Lake, Carley;McDonald, Katy;Parker, Kate;Sheridan, Danielle;Silver, Nick;Anderson, Brooke;Chang, Frank;Kansal, Sachin;Sullivan, Joe;	SEMANTIC
Primary Date	03/12/2016 11:33 pm	DOC_TYPE_ALIAS
Production Volume	JCCP_MDL106;	SEMANTIC
Redacted	No	SEMANTIC
Sort Date	01/28/2019 10:32 pm	SEMANTIC
SourceHash	0704653c760110e2fbad787b82a22764	SEMANTIC

Bouncer v3

Sunny Jeon
Data Scientist | Trust & Safety

March 28, 2016

Summary: This document provides an overview of *Bouncer*,[1] an intelligent decision system for anticipating and preventing safety incidents on the Uber platform (e.g., accidents, interpersonal conflicts). Statistical machine learning models predict whether drivers will or will not be involved in a safety incident 7, 14, 30, and 60 days into the future. High risk users are targeted for interventions that prevent safety incidents in randomized controlled trials. The key performance indicator is reduction in safety incidents.

Contents

[1] Previous versions of Bouncer: Bouncer v2+, Bouncer v2, Bouncer v1.

1. Motivation [back to contents]

As Uber usage continues to grow, the company may encounter a greater number and variety of safety risks—including fatal automobile accidents and interpersonal conflicts like physical altercations and sexual misconduct. Although safety incidents are rare and may appear random, **many safety incidents are predictable**. That is, safety incidents follow patterns and have precursors that can be leveraged to forecast them before they happen. For example:

- Safety incident rates increase dramatically between 1-3am, on Saturday nights/Sunday mornings, and on holidays and other days of major social gatherings (e.g., World Series, Outside Lands). [See analysis: Safety Hot Times/Days Analysis]
- Users that have caused business critical safety incidents have substantially more safety tickets going into the incident and substantially lower ratings in their first 10 trips than users that have never caused a major safety incident. [See analysis: Precursors to Business Critical Safety Incidents]
- Safety risks travel through social networks. That is, it is possible to predict a driver's safety ticket history using just their friends' ticket history. [See analysis: Safety Across Referral Networks]
- Safety incidents do not occur randomly across space, but are clustered in specific types of neighborhoods. [See analysis: Safety Hot Spots Analysis]

If safety incidents are predictable, they are preventable. This is the insight behind *Bouncer*, an intelligent decision system for anticipating and preventing safety incidents (e.g., accidents, interpersonal conflicts). Statistical machine learning models predict whether drivers will or will not be involved in a safety incident 7, 14, 30, and 60 days into the future. To prevent incidents from occurring, high risk users are targeted for a variety of interventions that have reduced safety incidents in randomized controlled trials. Key results from the project:

- At the **30/60 day forecasting windows**, predictive models correctly identify ~90% of drivers that cause dangerous driving incidents (with ~20% precision) and ~45% of drivers that cause interpersonal conflicts (with ~5% precision) in out-of-sample tests.
- At the **7/14 day forecasting windows**, predictive models correctly identify ~70% of drivers that cause dangerous driving incidents (with ~10% precision) and ~25% of drivers that cause interpersonal conflicts (with ~3% precision) in out-of-sample tests.
- Predictive models for both dangerous driving and interpersonal conflicts can be tuned to have >80% precision at the 30/60 day forecasting windows, but at the cost of recall.
- Interventions based on safety messaging **reduce safety incident counts by >10%** in controlled experiments conducted in the US (see SASSY Project Brief).

Bouncer is still very much a work in progress, but to solicit feedback and suggestions for improvement, this document describes Bouncer's incident reduction strategy, forecasting methodology, and results from performance evaluations.

Figure 1: Bouncer Intervention Framework

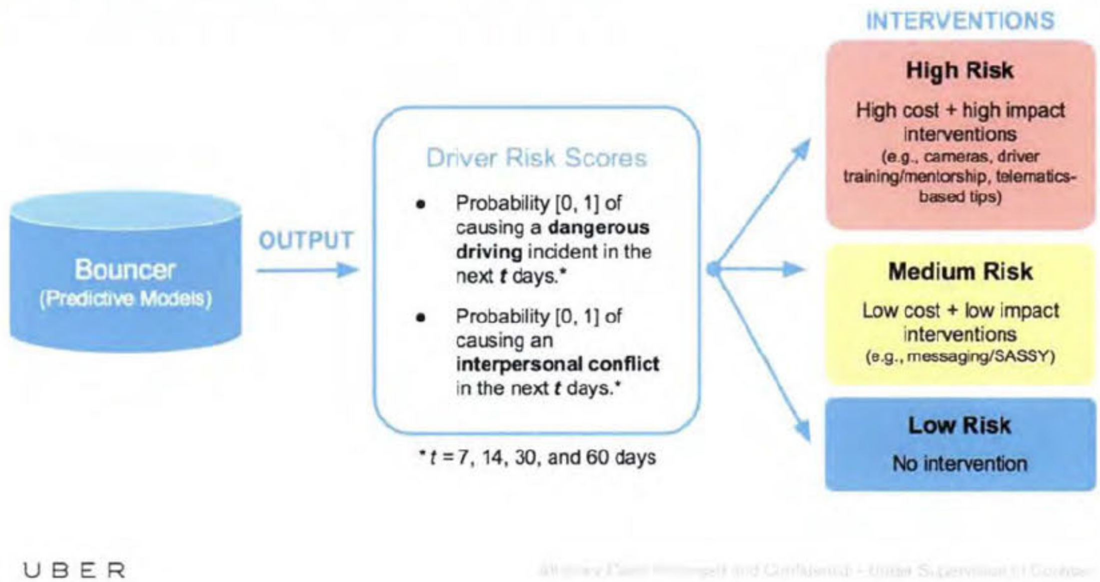

Figure 2: Products for Prevention by Stages of Trip Life Cycle

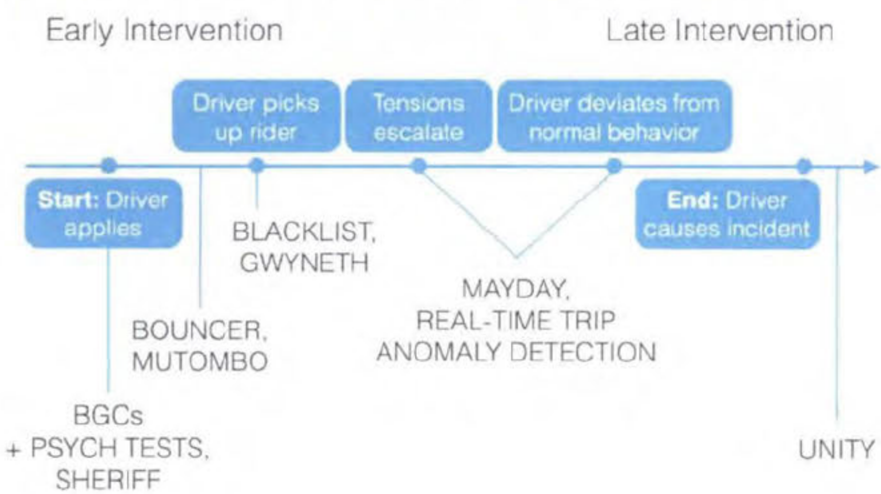

2. Safety Incident Prevention Strategy [back to contents]

Bouncer reduces safety incidents on the Uber platform by using patterns in historical data to identify which users are at greatest risk of causing a safety incident in the future. These users are then targeted for a variety of preventative interventions based on their specific risk profile. The incident reduction strategy is driven by 3 components:

1. Data Pipeline
2. Predictive Models
3. Automated Interventions

Figure 1 provides a high-level overview of the Bouncer intervention program. Figure 2 depicts where in the trip life cycle Bouncer -- as well as other safety products + initiatives -- intervene.

Data Pipeline

To compute risk scores, Bouncer draws on a data pipeline of internal and external data relevant for detecting high risk users. Data include:

[redacted]

The full variable list (including definitions and sources) is available here. In Section 5 -- Most Important Predictors, I identify the predictors with the greatest predictive power.

Predictive Models

Bouncer draws on the data pipeline to build a suite of statistical machine learning models for forecasting safety incidents. The current iteration of Bouncer (v3) focuses on identifying drivers/partners that are likely to cause two types of safety incidents **7, 14, 30, and 60 days into the future**:

- **Dangerous Driving:** accidents, distracted driving, poor/erratic driving, and traffic violations.
- **Interpersonal Conflicts:** verbal and physical altercations, inappropriate behavior, and sexual misconduct (instigated by the driver).

To accomplish this predictive task, the models leverage a variety of binary classification algorithms -- including Random Forests, Support Vector Machines, AdaBoost, and Stochastic Gradient Boosting -- to discriminate between drivers that are likely to cause a dangerous driving incident or interpersonal conflict, and those that are unlikely to. The best performing models are identified by systematically back-testing existing data, and predictive performance is estimated using an out-of-sample testing protocol.

Automated Interventions

Once Bouncer identifies drivers that are likely to cause a safety incident, the drivers are targeted for a series of interventions designed to prevent the incident from occurring. In other internal work, I have shown that the following safety messaging interventions are effective at reducing safety incident counts in randomized controlled trials:

- **Positive Reinforcement:** weekly messages acknowledging good behavior or good driving performance (see labs page).
  - Treatment reduces number of safety incidents by **-6%** ($p = 0.08$).
- **Rider Feedback:** weekly messages providing 3 randomly selected (scrubbed) verbatim comments from riders (see labs page).
  - Treatment reduces number of safety incidents by **-8%** ($p = 0.03$).
  - Results driven by impact on dangerous driving tickets, which saw a **-10%** decrease compared to control group ($p < 0.01$).
- **SASSY (Smart + Automated Safety SMS System):** personalized safety messages that are algorithmically composed based on each driver's weekly performance (see SASSY Project Brief for more details).
  - SASSY -- and in particular the reactive bundle of treatments (pro-tips + warnings for repeat violations) -- reduces safety incident counts by **-12%** ($p < 0.01$) and dangerous driving ticket counts by **-14%** ($p < 0.01$).

These messaging interventions have intentionally been designed to be subtle and low-cost because the interventions are targeted towards those that are *predicted* to cause a safety incident -- not necessarily those that have already caused an incident. As such, strong interventions like banning/off-boarding are probably inappropriate. However, if the predictive models can demonstrate reliability over time, it may be worth experimenting with moderately costly but potentially high-impact interventions, like restricting use during high-risk times (1-3am), requiring driver training, or installing cameras in partner

vehicles for audio, video, GPS, and/or telematics monitoring. This slide deck describes the prevention action playbook.

Figure 3: Bouncer Testing Process

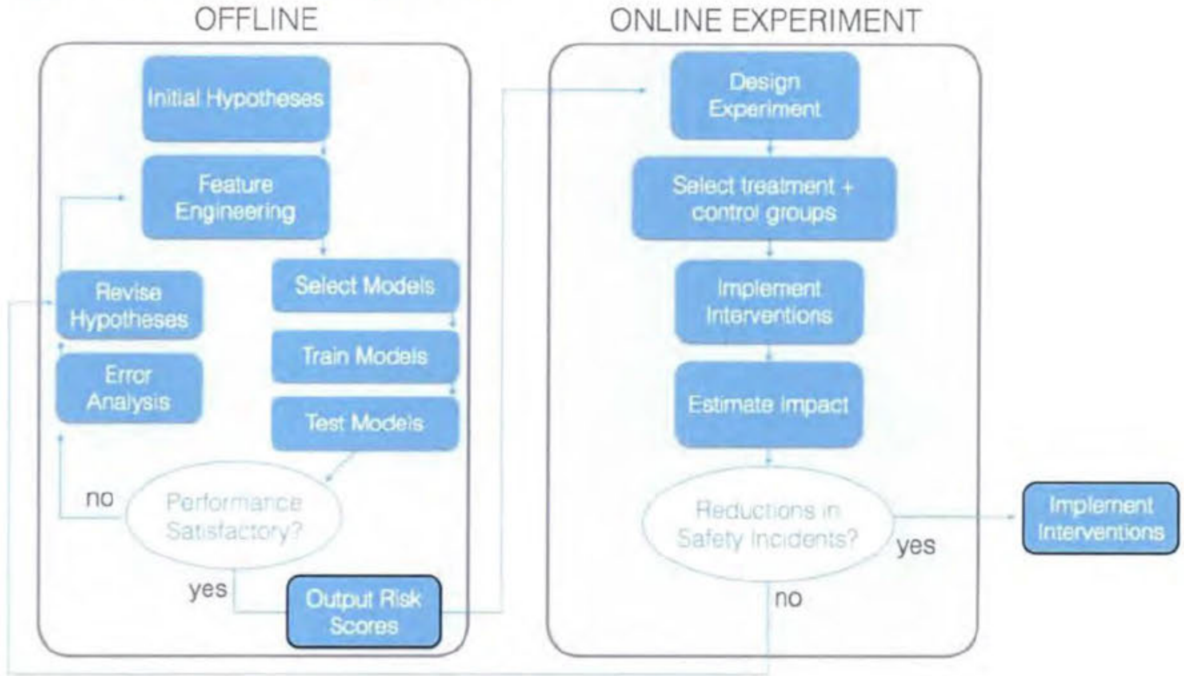

Testing Process

Figure 3 depicts the process by which Bouncer is developed and tested. The process begins with offline experimentation, which involves an iterative process of building predictive models, testing them, and revising them based on the test results and findings from error analysis. Once the predictive performance of the models is satisfactory, Bouncer is launched in select cities using a randomized experimental framework that randomly divides drivers into treatment (receive Bouncer interventions) and control (do not receive Bouncer interventions) groups to isolate and estimate the impact of Bouncer on safety incident counts. If the experiment results indicate Bouncer reduces incident counts, it is scaled up to other cities. If the experiment results cannot identify effects from Bouncer, both the interventions and the predictive models are re-designed by repeating the entire offline-online experiment process.

[redacted]

Data

The current iteration of Bouncer focuses on detecting high risk drivers in 4 major US markets: San Francisco, Los Angeles, Chicago, and New York City.

To build models for these cities, I collect data on a stratified random sample of 28,986 drivers (6,805 Chicago, 7,517 LA, 7,538 NYC, 7,126 SF). Because safety incidents are uncommon, users with safety incidents are over-sampled and users without safety incidents are under-sampled to balance the data (this dataset is further balanced before and during model training using simple down-sampling, internal down-sampling, and SMOTE, see Figure 4 for details). This sample was drawn by pulling:

- All 705 SF, LA, Chicago, and NYC drivers that have received at least one L3 and L4 tickets between January 1 - September 30, 2015,
- All 4,151 SF, LA, Chicago, and NYC drivers that have received at least one L2 ticket between January 1 - September 30, 2015,
- A 10% random sample of SF, LA, Chicago, and NYC drivers that have received at least one L1 ticket between January 1 - September 30, 2015 (12,134 drivers), and
- A random sample of 11,996 SF, LA, Chicago, and NYC drivers.

For each user, I create a day-level time-series that goes from January 1, 2015 to September 30, 2015, and includes data on a variety of user features built from Zendesk ticket data and user account data (e.g., changes in ratings, ETAs, cancellation rates, trip counts, times worked, fares earned/paid, etc.). These data are operationalized as over 200 distinct predictors. This Variable List provides definitions and sources.

Data Pre-Processing

The data are pre-processed in several ways:

- All predictors are centered and scaled to have mean 0 and standard deviation 1. In the model training process, this normalization is made within resampling loops.
- All predictors with near zero variance are removed. Near zero variance predictors are those that meet two conditions: (i) less than 10% of observations are unique, and (ii) the ratio of frequencies for the most common value over the second most common value is greater than 95/5.
- Highly correlated predictors are removed by computing pair-wise correlations, identifying highly correlated predictors (>0.80), and then removing the predictor with the higher mean pair-wise correlation across the full correlation matrix.

[redacted]

[redacted]

The purpose of the probe set is to optimize additional model parameters and specifications to fit the model for our purposes, such as the the optimal subsampling method for addressing the rare events problem (simple down-sampling v. internal down-sampling v. SMOTE), and to identify the optimal probability cutoff for maximizing recall without sacrificing precision. The following is a full list of the different specifications I test:

- Subsampling techniques:
  - Simple down-sampling
  - Internal down-sampling
  - SMOTE

- Optimal probability cutoff for maximizing recall without sacrificing precision

- Operationalization of dependent variable
  - Safety incident count ≥ 1 over next 7, 14, 30, 60 days
  - Safety incident count ≥ 2 over next 7, 14, 30, 60 days
  - **Note*: Bouncer v2 predicts incident counts using raw Zendesk ticket data
  - **Note*: Bouncer v1 predicts incident rates (incident count/trip count)

- Subset
  - Subset to drivers that got activated during the sample time period
  - Data aggregated by city
  - Data aggregated across all cities
  - Remove observations where incident count = 1 on the dependent variable (to reduce noise during model training process)

- Predictors
  - Full set of predictors
  - Top 50 predictors (by variable importance scores)

After model specifications are optimized against the probe set, a separate random sample of 17,836 drivers from the same US cities (4,398 Chicago, 4,572 LA, 4,484 NYC, 4,382 SF) is selected for the ***test set***. The accuracy of predictions for the random test set is used to assess the predictive performance of the final models -- which is what we would expect to see "in the wild" if Bouncer were launched.

4. Predictive Performance in US Markets [back to contents]

Performance Metrics

The primary performance metrics are precision, recall, and percent positive predictions on the out-of-sample random test set, where:

$$Precision = \frac{True\ Positives}{True\ Positives\ +\ False\ Positives}$$

$$Recall = \frac{True\ Positives}{True\ Positives\ +\ False\ Negatives}$$

$$Percent\ Positive\ = \frac{True\ Positives\ +\ False\ Positives}{Total\ Predictions\ Made}$$

Because the costs of safety incidents are high, a good model will have high recall to minimize false negatives. An ideal model, however, will also have high precision, particularly when the costs of false positives are nontrivial. Percent positive gives the percent of "positive" (high-risk) predictions. Because one prediction is generated for every driver at every point in time, percent positive gives the percent of the entire driver population that needs to be targeted to reach recall. Larger recall:percent positive ratios thus indicate stronger performance.

Main Results

Bouncer can be optimized for many different purposes and performance metrics. Suppose, for example, that our sole interest is in identifying 100% of drivers that are likely to cause a dangerous driving incident or interpersonal conflict 30 days into the future (to be as risk-averse as possible). Then, the probability cutoff used for classification can be manipulated to maximize recall at the cost of precision. Alternatively, suppose we want to maximize both recall and precision due to some cost produced by false positives.[2] Then the probability cutoff can be manipulated to maximize the F_1 *Score*, which is the harmonic mean of recall and precision (F_1 *Score* = 2*precision*recall/(precision + recall)). To place more weight on recall, the F_2 *Score* and F_3 *Score* are options. The generic formula for computing an F_β Score is:

[2] In the case of Bouncer, the cost of false positives is the unnecessary spamming of safety messages to our drivers/partners. Given the volume of messages that are already sent to our drivers/partners by ops, comms, supply, and many other Uber teams, there is an interest in reducing unnecessary spamming, which would likely diminish the value of each individual message that is sent to drivers/partners (see SMS Volume Report). It would also likely further reduce the already low email open rate (anecdotally at ~30%), and potentially reduce the efficacy of each individual message.

[redacted]

Figure 6: Interpersonal Conflict Predictions

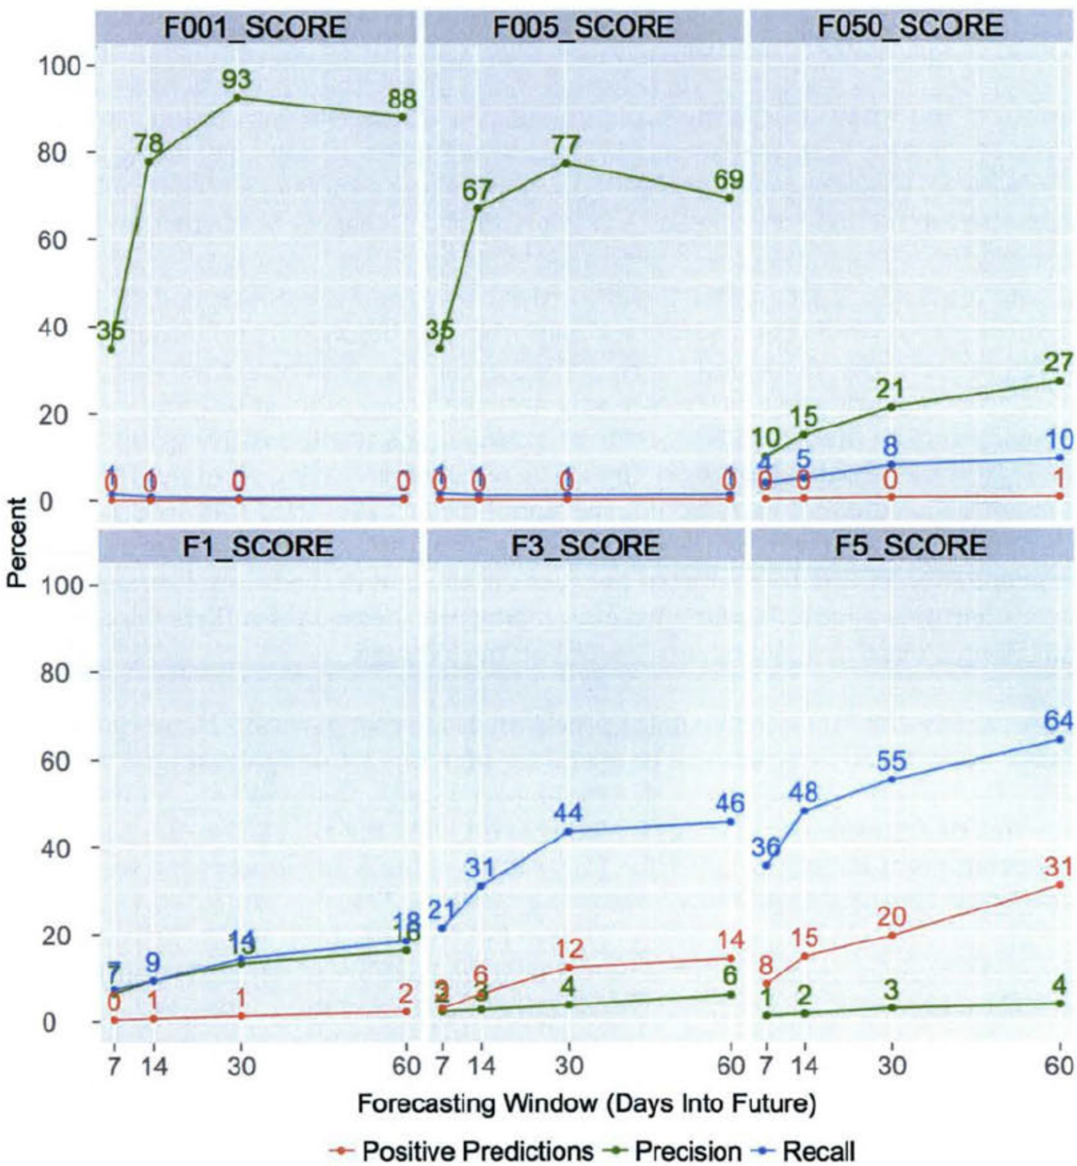

To demonstrate these different applications of Bouncer, Figures 5-6 present the precision, recall, and percent positive for out-of-sample predictions when selecting the probability cutoffs that maximize the $F_{0.01}$ Score, $F_{0.05}$ Score, $F_{0.50}$ Score, F_1 Score , F_3 Score, and F_5 Score.

[redacted]

3. In fact, it correctly anticipates nearly 100% of drivers with ≥3 tickets in the SF, LA, Chicago, and NYC test sets. These results are favorable, as it implies the model is good at anticipating the highest-risk drivers -- the ones we should care most about.

- Results for the driver interpersonal conflict models are not as strong.

[redacted]

[redacted]

[redacted]

Comparison to Naive Forecasting Strategies

The predictive power of Bouncer can be further validated by comparing them against benchmarks set by naïve forecasting strategies that use simple heuristics and no model, such as:

[redacted]

[redacted]

5. Improvements From Previous Models [back to contents]

[redacted]

[3] See Davis and Goadrich (2006) at http://www.autonlab.org/icml_documents/camera-ready/030_The_Relationship_Bet.pdf.

[redacted]

6. Most Important Predictors [back to contents]

[redacted]

[redacted]

[redacted]

Table 3: Bouncer Development and Roll-Out Schedule

	Q4 2015	Q1 2016	Q2 2016	Q3 2016
VERSION	v2	v3	v4	v5
TARGET	• Driver-instigated incidents			• Driver-instigated incidents • Rider-instigated incidents
DATA PIPELINE	• Trips • Safety tickets • Driver background	• Telematics • Weather • BGCs • Text analysis on feedback	• NLP on feedback • Hot spots • Referral networks • Event schedules • External	
COVERAGE	SF, LA, CHI, NYC	SF, LA, CHI, NYC	All US Cities	US, Europe, ANZ
ENGINEERING	v2 data pipeline	v2 data pipeline	v3 data pipeline	v4 data pipeline
KPIs	• Safety incidents (counts + rates) • Lawsuits / settlements (number + cost) • Predictive performance (precision + recall)			

7. Challenges [back to contents]

- Data integrity — lots of measurement error in safety data (e.g., Zendesk ticket classifications).
- Safety incidents are rare events (e.g., <0.10% of observations), creating highly imbalanced data.
- Model training process very slow despite leveraging parallel processing.

8. Next Steps [back to contents]

- Productionize predictive models using Michelangelo.
- Error analysis. Figure out why some high risk users are being incorrectly classified as low risk and why some low risk users are being incorrectly classified as high risk. Then build new features to capture this variation.
- Maximize precision while holding recall as close to 100% as possible. Strategies:
  - Increase sample size.
  - Build better features that distinguish safe users from unsafe users.
  - Better tuning parameters.

  - In model training, increase costs for misclassification of positive cases (i.e., cost sensitive learning).
- Build features for Bouncer v4 (features listed on this variable list).
- Develop models for other cities outside of US. Need to figure out if there's variation in predictive performance across cities, and if so, what explains the variation (e.g., young v. mature market, regional/cultural differences, etc.)
- Build models for predicting **rider instigated interpersonal conflicts** -- one of the most common causes of L3 and L4 safety incidents.
- Figure out how best to operationalize Bouncer's predictions. Lots of potential interventions. Need to test them using randomized experiments to identify the most effective strategies.
- Table 3 describes the Bouncer development and roll-out plan.

9. Team [back to contents]

Data: Sunny Jeon, Spencer Boucher, David Purdy
Engineering: Rami Mawas, Yisheng Liang, Ron Tal, Jeremy Hermann
Safety Ops: Jeana Williams, Becky Mar
Product: Dhruv Tyagi, Dima Kovalev
Legal: Justin Suhr, Seth Schreiberg

Appendix [back to contents]

Figure A.1: Predicted Probability versus Actual Incident Count

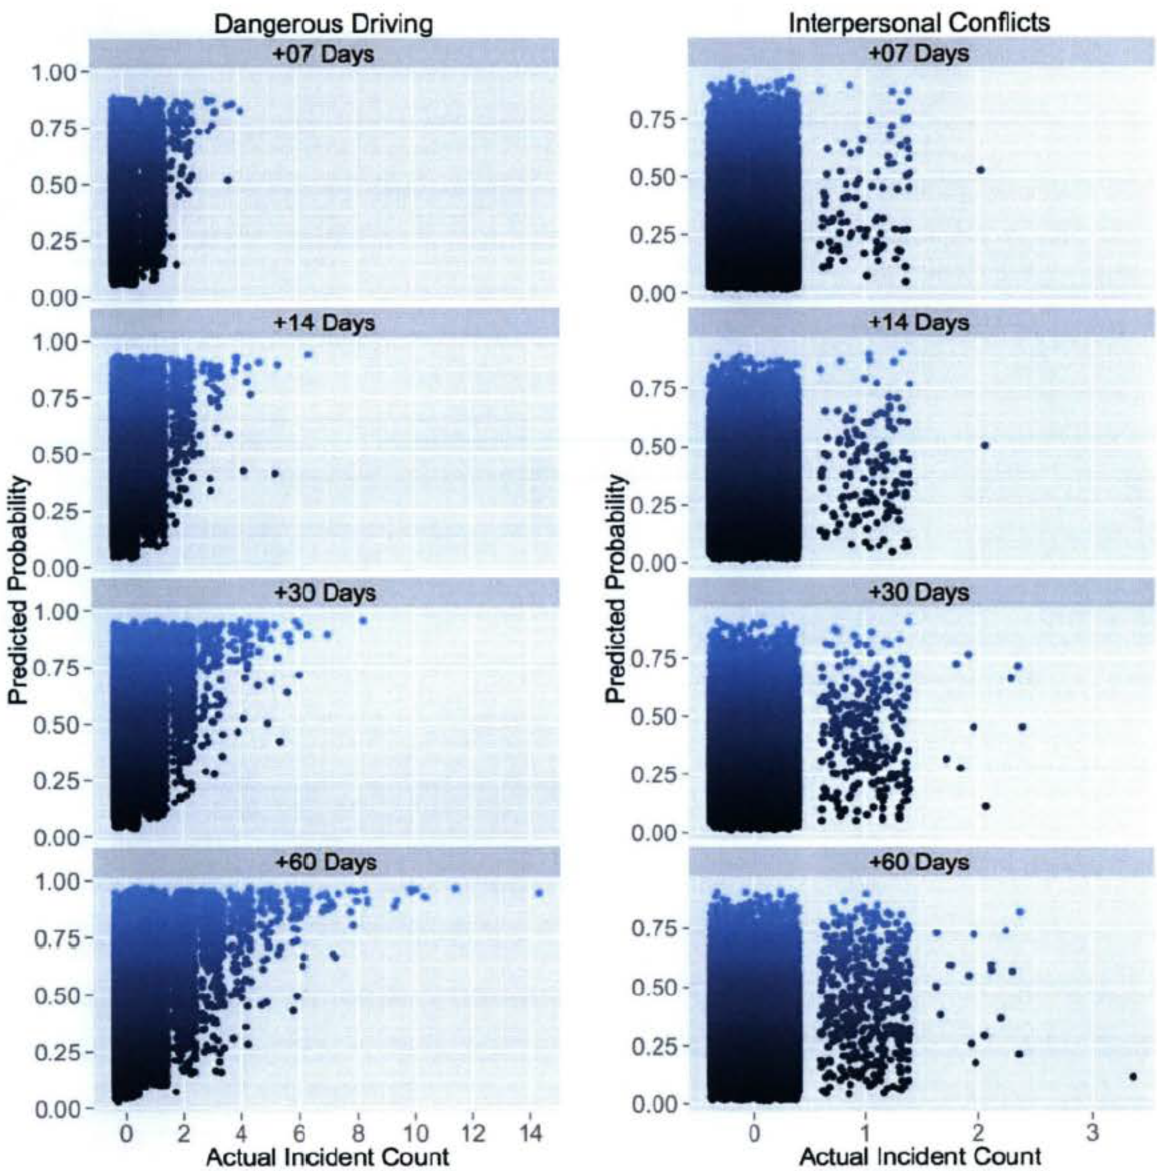

Figure A.2: Comparing Model Performance (Test Set AUROC)